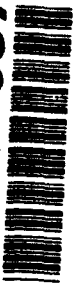


AD-A279 697



Fiscal Year Report

Principal Investigators: Drew McDermott and Gregory Hager
Institution: Yale University Department of Computer Science
Phone: (203) 432-1997
E-mail: mcdermott@cs.yale.edu
Title: Knowledge-Based Planning
Grant Number: N00014-91-J-1577
Reporting Period: 1 Oct 1992 - 30 Sep 1993

DTIC
ELECTE
MAY 31 1994
S F D

1 Productivity Measures

- Refereed papers submitted but not yet published: 3
- Refereed papers published: 3
- Unrefereed reports and articles: 2
- Books or parts thereof submitted but not yet published: 0
- Books or parts thereof published: 0
- Patents filed but not yet granted: 0
- Patents granted (include software copyrights; provide patent numbers): 0
- Invited presentations: 7
- Contributed presentations: 5
- Honors received: 1 (Drew McDermott served as co-editor of *Artificial Intelligence Special Issue on Planning*, to appear next year)
- Prizes or awards received: 0
- Promotions obtained: 0
- Graduate students supported $\geq 25\%$ of full time: 3
- Post-docs supported $\geq 25\%$ of full time: 1
- Minorities supported: 1

This document has been approved
for public release and sale; its
distribution is unlimited.

DTIC QUALITY INSPECTED 1

94 5 27 108

94-16142



467 051

~~93 11 24 119~~

Fiscal Year Report

Principal Investigators: Drew McDermott and Gregory Hager
Institution: Yale University Department of Computer Science
Phone: (203) 432-1997
E-mail: mcdermott@cs.yale.edu
Title: Knowledge-Based Planning
Grant Number: N00014-91-J-1577
Reporting Period: 1 Oct 1992 - 30 Sep 1993

2 Detailed Summary of Technical Progress

Our work has focused on planning and execution for agents with imperfect sensors. Our focus has been on models that tightly couple sensing to action, but we have also studied models of high-level planning in simulated dynamic worlds, where the emphasis is on projecting and analyzing plans.

Our overall model of planning is this: An agent must be continuously executing plans in order to make progress on its goals. These plans are driven by sensors, and can normally cope with deviations from expected results without intervention from the planner. When intervention is required, the planner starts from scratch, generating plans and then revising them on the basis of their projected results. The projector contains a probabilistic model of the world that allows the planner to forecast probable errors.

Our work on sensor-driven action has focused on these areas:

1. Theoretical foundations for set-based decision-making algorithms.
2. Visual tracking and vision-based control of servo systems.
3. Comparison of set-based and statistically-based estimation.

Our algorithms for set-based decision-making explicitly recover the parameters of geometrical or physical models until decision-specific accuracy criteria have been met. It is possible to supply any number of decision criteria that are evaluated in parallel as parameter recovery is performed. We have been able to show that the algorithm we use for evaluating decision criteria is correct and complete except for a vanishingly small set of problems. Correctness means that only physically correct decisions will be made as long as the input data are consistent with the parametric models supplied to the algorithm. Completeness means that the algorithm will terminate on all inputs. We have shown that the algorithm we use will terminate on all inputs except for certain boundary cases that are typically a set of measure zero in the space of algorithm inputs.

These results hold for an extremely wide variety of problem settings including problems where the number and type of geometric models is not known a priori.

Accession For	
NTIS	CRA&I <input checked="" type="checkbox"/>
DTIC	TAE <input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By <i>per lti</i>	
Distribution /	
Availability /	
Dist	Availability / Special
A-1	

This means that our algorithm is effectively a decision procedure even when both segmentation and parameter fitting must be performed simultaneously.

Similar set-based methods have been employed by us for working in mobile robot mapping and navigation. At the same time, several researchers have used classical statistical estimation methods for the same purposes. This year we have undertaken a study to compare statistical estimation methods with set-based methods for the purpose of robot navigation. We have found that set-based methods typically outperform statistical methods when the estimation problem has low-dimension and is nonlinear. As it turns out, many of the estimation problems faced in robot navigation have these properties, so set-based methods would be expected to outperform statistical methods. We have run both simulated and real experiments with a mobile robot system and have verified that this is true. These experiments were done in the context of a map-building system that uses interval-based methods to represent its current model of its surroundings, and to correct the model as new information comes in.

Our final effort in the area of sensor-driven plans is the implementation of a system for real-time tracking in images using only standard framegrabber hardware and standard workstations. This system, written in C++, provides a variety of tracking "templates" that permit a user to define a collection of tracked objects and to enforce constraints among them. For example, a basic trackable feature may be contrast edges in an image. A corner can then be defined as the intersection point of two contrast edges, the face of a rectangle as a plane containing four corners, and a box as a collection of six planes. Each of these concepts, a corner, a face, and a box, can be defined in terms of the preceding elements, and have particular state variables associated with them. So, for example, a corner has three variables (two positions and one orientation), a plane may have two (a slant and a tilt) or more, and a box may have three to six as the user desires. This tracking system is currently being used to implement a visual servoing system for mobile robot navigation and manipulation of objects.

In high-level planning, our work has been concentrated in these areas, all in the context of our transformational planner XFRM:

1. Development and experimental testing of an architecture for interleaving planning with execution.
2. Further exploration of methods for representing executable plans declaratively.
3. Optimization and verification of a probabilistic time map.

As reported last year, we have implemented a robust interface between a planner and an execution module that allows swapping in of a new plan at any time. We have run preliminary experiments in our "delivery world" that show that the system is able to achieve significant improvements in performance times simply by planning simultaneously with execution. Typically, in cases where

the planner can run fast enough to "beat" the interpreter, the plan it swaps in embodies speedups over the default plan that compensate for the time and side effects incurred while executing the default.

Transforming plans is a difficult business. To make it easier, plans have to be represented in a transparent manner that enables the planner to see the purposes of the pieces of behavior it encapsulates. We have added more declarative constructs for expressing these purposes. In particular, where possible we express sensor tests using a BELIEF construct that makes it explicit which beliefs about the agent's environment are being polled. We have also developed a Prolog-like "meta-language" XFRM-ML for expressing transformation rules and their associated preconditions.

We have completely rewritten our probabilistic projector so that it runs about twenty times faster and embodies more functionality. The projector now includes rules that model autonomous Poisson-distributed events. The rule $(P \rightarrow E \ P \ d \ E)$ specifies that over any interval where P is true, the expected time to the next occurrence of E is d . We have developed a formal semantics for these and the other rule types, and shown that the program generates timelines with the probabilities given by the formal semantics.

Fiscal Year Report

Principal Investigators: Drew McDermott and Gregory Hager

Institution: Yale University Department of Computer Science

Phone: (203) 432-1997

E-mail: mcdermott@cs.yale.edu

Title: Knowledge-Based Planning

Grant Number: N00014-91-J-1577

Reporting Period: 1 Oct 1992 - 30 Sep 1993

3 Publications, Presentations and Reports

- Sami Atiya and Gregory D. Hager, "Real-Time Vision-Based Robot Localization," *IEEE Trans. on Robotics and Automation*, in press
- Michael Beetz, "Improving Robot Plans During their Execution," presentation at the Bavarian Center for Knowledge-based Systems/ University of Erlangen; the Technical University of Darmstadt; and the German Research Center for Artificial Intelligence, Saarbruecken, Germany, August, 1993
- Sean Engelson, "Active Place Recognition Using Image Signatures." SPIE 1992
- Sean Engelson and Niklas Bertani, "ARS MAGNA: The Abstract Robot Simulator Manual." Yale Computer Science Report 928, 1992
- Sean Engelson and Drew McDermott, "Maps Considered As Adaptive Planning Resources". AAAI Fall Symposium on Applications of AI to Real-World Autonomous Mobile Robots, 1992
- Greg Hager, "Towards task-directed planning of cooperating sensors," presentation at the SPIE Sensor Fusion Workshop, 1992
- Greg Hager, "Sensor planning for reactive robot programs," invited presentation at the Allerton Conference on Control and Computing, 1992
- Greg Hager: "Efficient Solution of Large Systems of Nonlinear Constraints With Inexact Data and Explicit Termination Criteria," presentation at the Conference on Numerical Computation with Automatic Result Verification in LaFayette Louisiana, Feb 25-Mar 1.
- Gregory Hager, "Task-Directed Computation of Qualitative Decisions from Sensor Data," presentation at IEEE Conf. on Robotics and Automation, May, 1993

- Gregory D. Hager, "Solving Large Systems of Nonlinear Constraints with Application to Data Modeling," *Interval Computations*, in press
- Gregory D. Hager, "Task-Directed Computation of Qualitative Decisions from Sensor Data," submitted to the *IEEE Trans. on Robotics and Automation*
- Greg Hager, Sean Engelson, and S. Atiya. "On Comparing Statistical and Set-Based Methods in Sensor Data Fusion (with S. Engelson and S. Atiya)." *Proc. International Conference on Robotics and Automation*, 1993
- Drew McDermott, "Transformational planning of reactive behavior." Yale Computer Science Report 941, December, 1992.
- Drew McDermott: "Transformational Planning of Reactive Behavior," presentation at the MIT Workshop on Autonomous Underwater Vehicles, Cambridge, MA, January 26-27, 1993
- Drew McDermott, "Planning is Automatic Programming for Agents," invited talk at the AAAI Spring Symposium, March 22-24, 1993
- Drew McDermott, "Transformational Planning of Reactive Behavior," Brandeis University, April 8, 1993
- Drew McDermott, panel discussion at AAAI on "Pros and Cons of Software Evaluation," July, 1993

Fiscal Year Report

Principal Investigators: Drew McDermott and Gregory Hager
Institution: Yale University Department of Computer Science
Phone: (203) 432-1997
E-mail: mcdermott@cs.yale.edu
Title: Knowledge-Based Planning
Grant Number: N00014-91-J-1577
Reporting Period: 1 Oct 1992 - 30 Sep 1993

4 Transitions and DoD Interactions

Prof. McDermott has been serving on the Technical Review Board for the DARPA Transportation and Scheduling Initiative. The purpose of the board is to provide high-level feedback to researchers in this area, using insights gained from past research on planning and scheduling.

Fiscal Year Report

Principal Investigators: Drew McDermott and Gregory Hager

Institution: Yale University Department of Computer Science

Phone: (203) 432-1997

E-mail: mcdermott@cs.yale.edu

Title: Knowledge-Based Planning

Grant Number: N00014-91-J-1577

Reporting Period: 1 Oct 1992 - 30 Sep 1993

5 Software and Hardware Prototypes

1. We are still supporting the Reactive Plan Language interpreter (in Common Lisp), and the constraint solver (in C and C++).
2. We are sharing tracking and mobile-robot control software with interested institutions.
3. The probabilistic time map has been split off from the planner, and is now available via anonymous ftp.

Fiscal Year Report

Principal Investigators: Drew McDermott and Gregory Hager

Institution: Yale University Department of Computer Science

Phone: (203) 432-1997

E-mail: mcdermott@cs.yale.edu

Title: Knowledge-Based Planning

Grant Number: N00014-91-J-1577

Reporting Period: 1 Oct 1992 - 30 Sep 1993

6 Photographs, Vugrafs and Videotapes

Agent Architecture